arXiv:1606.04593v1 [cs.CR] 14 Jun 2016

Multi-Party Encrypted Messaging Protocol design document

Release 0.3-2-ge104946

Mega Limited, Auckland, New Zealand
Guy Kloss <gk@mega.co.nz>

11 January 2016

1 Strongvelope Multi-Party Encrypted Messaging Protocol **2**
1.1 Intent . . . 2
1.2 Security Properties . . . 2
1.3 Scope and Limitations . . . 3
1.4 Assumptions . . . 3
1.5 Messages . . . 3
1.6 Terminology . . . 4

2 Cryptographic Primitives **5**
2.1 Key Pair for Authentication . . . 5
2.2 Key Agreement . . . 5
2.3 (Sender) Key Encryption . . . 5
2.4 Message Authentication . . . 6
2.5 Message Encryption . . . 6

3 Message Encryption **7**
3.1 Sender Key Exchange . . . 7
3.2 Content Encryption . . . 8

4 Protocol Encoding **9**
4.1 TLV Types . . . 9
4.2 Message Signatures . . . 10
4.3 Message Types . . . 10
4.4 Keyed Messages . . . 10
4.5 Followup Messages . . . 10
4.6 Alter Participant Messages . . . 11
4.7 Legacy Sender Key Encryption . . . 11
4.8 Sender Keys to Include . . . 11

5 Requirements for Messaging Workflows **12**
5.1 Initialising Chat Room Participation . . . 12
5.2 Update Sender Key . . . 12
5.3 Seed Encryption Handler . . . 12
5.4 Rotate Key . . . 13
5.5 Key Reminders . . . 13
5.6 Alter Participant Messages . . . 13

6 Strongvelope Evolution **14**

Bibliography **15**

© 2015 Mega Limited, Auckland, New Zealand.
https://mega.nz/

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License. For the full text of the license in various formats, and other details, see https://creativecommons.org/licenses/by-sa/4.0/

To give a non-legal human-readable summary of (and *not* a substitute for) the license, you are free to:

- Share – copy and redistribute the material in any medium or format
- Adapt – remix, transform, and build upon the material

for any purpose, even commercially. The licensor cannot revoke these freedoms as long as you follow the license terms:

- Attribution – You must give appropriate credit, provide a link to the license, and indicate if changes were made; but not in any way that suggests the licensor endorses you or your use.
- ShareAlike – If you distribute your adaptions of the material, you must do this under the same license as the original. You may not apply additional restrictions when doing so, legal or technical.

No warranties are given. Other legal rights may extend (e.g. fair use and fair dealing) or restrict (e.g. publicity, privacy, or moral rights) your permission to use this material, outside of the freedoms given by this license.

CHAPTER

ONE

STRONGVELOPE MULTI-PARTY ENCRYPTED MESSAGING PROTOCOL

In this document we describe the design of a multi-party messaging encryption protocol “Strongvelope”. We hope that it will prove useful to people interested in understanding the inner workings of this protocol as well as cryptography and security experts to review the underlying concepts and assumptions.

In this design paper we are outlining the perspective of chat message protection through the Strongvelope module. This is different from the *product* (the Mega chat) and the transport means which it will be used with. Aspects of the chat product and transport are only referred to where appropriate, but are not subject to discussion in this document.

1.1 Intent

The Strongvelope protocol is intended to protect the privacy and confidentiality of content exchanged in Mega’s chat application among groups of participants.

As underlying design goals a set of security properties needs to be met (see next section). These have to be upheld in instant as well as offline messaging. Furthermore, upon changes of the set of chat participants, it needs to be ensured that the used encryption keys change. It is required that newly added participants are unable to decrypt previously sent messages, and former (now excluded) participants are unable to decrypt messages after their exclusion.

1.2 Security Properties

Our design rationale was based on the need of implementing a multi-party capable instant messaging protocol, that would be light enough to be executed over the Internet/World Wide Web when implemented in a JavaScript text chat client or on mobile devices. The protocol would need to provide primarily confidentiality of the conversation as well as message and message origin authenticity.

- **Confidentiality**
  of the conversation, so its content is not accessible or readable by an outsider.
- **Message authenticity**
  against alteration during transport or storage. A recipient can be assured that the message is arriving exactly as it has been originally authored.
- **Message origin authenticity**
  against both outsider intrusion and the impersonation of existing participants by other malicious participants in the session. This means that the user can be assured of the authenticity of the sender of each original message even if other participants in the room try to impersonate the sender and send messages on their behalf.

1.3 Scope and Limitations

Any further assurances underpinned by cryptography are not (yet) supported by this protocol. Further security enhancing properties may be introduced in the future evolution of Strongvelope or are part of a different encryption module featuring a different (but overlapping) set of properties (such as the mpENC protocol *[mpENC]*).

1.4 Assumptions

The design of Strongvelope was driven by a number of constraining assumptions.

1.4.1 Underlying Transport

Strongvelope is intended to be a protocol design that is agnostic of the used transport protocol. However, we are presuming that the underlying transport mechanism message will retain message order.

1.4.2 Chat "Rooms"/Transport Channels

Group chats require a shared transport channel, often referred to as a "room" in which the participants of the chat gather and receive messages from the group. We are assuming that there is at the most *one* Strongvelope multi-party chat present per transport channel (room). Not all members of the channel need to be participants in the Strongvelope chat session. This is intended, as members may (more or less freely) be added to or removed from participation in a channel, but they explicitly need to be added to, or removed from the Strongvelope participants in a session. Therefore it is expected that there are potentially a few extra members in the chat room, which are (not or not yet) participating in the encrypted session.

Note that in this context a "chat room" or a "channel" are synonymous for the underlying technical message exchange mechanism used for the group communication. In contrast a "chat" is the actual exchange of messages on top of the room or channel.

1.4.3 Transport Meta-Data

It is assumed that the transport mechanism will carry a limited amount of meta-data for message delivery, which is to be used for the encryption protocol to work. This includes specifically a sender (originator) ID for each message as well as a unique ordering criteria (enforced by the server) to place all messages of a single chat room within a total order as seen by the server.

1.4.4 Implementation

To make the implementation robust for working on many types of transports, some precaution should be taken when processing incoming protocol messages. For example, due to the broadcast nature of some protocols (e.g. multi-party chat rooms or IRC channels), there may not be the possibility for directed message delivery, therefore an implementation may adopt the option to abort processing of messages not intended for the client itself. Specifically, these may be messages not intended for oneself sent over a broadcast channel (e.g. a message sent by oneself or a message received after quitting participation while still a member of the chat "room" on the transport implementation).

1.5 Messages

It is not intended for Strongvelope to co-exist with clear-text messaging. It is assumed that *all* messages will be protected consistently via the Strongvelope protocol. All messages contain content required to drive the crypto-

graphic protection, and they may also contain an encrypted data payload. All content in Strongvelope messages (besides the first byte, indicating the protocol version used) is cryptographically signed.

1.6 Terminology

In this document the term “message” is used frequently. To avoid ambiguity, a “message” is considered to be a unit of information transported by means of any suitable wire protocol. Such a message includes additional information and meta-data, such as cryptographic keys, signatures, etc. A user content message in the general sense is authored by a user/participant of the chat, and is termed to be a “message payload” or just “payload”. Therefore it is possible to send a message without a payload, which is also termed to be a “blind message”, as it will not need to be displayed within a client.

CHAPTER

TWO

CRYPTOGRAPHIC PRIMITIVES

Mega provides a cloud-based platform enjoying a large popularity. Therefore, server-side scalability is of importance as well as feasibility for the messaging concept and the client implementation. Even though server scalability tends to be orthogonal to messaging encryption protocols, they have shown to add a significant overhead on the server infrastructure (due to the number of "blind" protocol bootstrapping messages and increased message size), so that fewer users can be served per server provided.

Out of experience, most of the clients will be using the messenger on the Mega platform through a Web or mobile client with limited computational capabilities (e.g. with end-to-end cryptography implemented in JavaScript). To maximise user experience through a fast response time (less lag) and reduced load on the executing end point hardware (which often are mobile devices), it is desirable to avoid frequent "heavy" computing operations (e.g. frequent exponentiation of big integers).

The cryptographic primitives have generally been chosen to match a general security level of 128 bit of entropy on symmetric ciphers such as AES. This is equivalent to 256 bit key strength on elliptic curve public-key ciphers and 3072 bit key strength on discrete logarithm problem based public-key ciphers (e.g. RSA, DSA). However, 3072 bit key strength is considered to be too expensive in many cases (computationally as well as with its demand on the entropy source). For security reasons NIST standardised ECC curves have been avoided*[NIST-ECC-Failure]*.

2.1 Key Pair for Authentication

Authentication of the sender as well as the message transmitted is performed via a cryptographic signature. For computational efficiency and good implementation across the different environments a signature key pair using EdDSA with the Edwards curve 25519 ("Ed25519") *[Ed25519]* is used. It is comparably compact (256 bit public keys, 512 bit signature size, 256 bit private key seed size).

2.2 Key Agreement

To agree on a symmetric encryption key between a sender and recipient a Diffie-Hellman (DH) scheme is used. To be more easily suitable for offline messaging capability, static "chat keys" using the Montgomery curve 25519 ("Curve25519") *[Curve25519]* are used. Pair-wise encryption keys (between sender and receiver) are derived using ECDH (elliptic curve Diffie-Hellman).

2.3 (Sender) Key Encryption

Each sender's key is encrypted with 128 bit AES in cipher-block-chaining (CBC) mode. No padding is used (all encryption keys a sender is using are exactly the size of a block).

2.4 Message Authentication

Each message is signed cryptographically for sender as well as message authenticity. To avoid message inflation and computational strain on the clients, an elliptic curve (EC) signature scheme is used. For efficiency, security and reputation the Edwards curve 25519 (“Ed25519”) *[Ed25519]* has been chosen (256 bit public keys, 512 bit signature size, 256 bit private key seed size) providing EdDSA signatures.

2.5 Message Encryption

Messages are encrypted with 128 bit AES in counter (CTR) mode.

CHAPTER

THREE

MESSAGE ENCRYPTION

Every sender is responsible for generating their own symmetric encryption key, ensuring user-controlled keys for any encrypted content sent. Therefore, each sender generates their own symmetric encryption keys used for encrypting the message payload. These "sender keys" then need to be exchanged with all other participants within a chat (pair-wise). The sender keys as well the message payload content then need to be encrypted for message transport and storage.

3.1 Sender Key Exchange

For this purpose "keyed" messages exist, which will carry the sender key to all recipients. Each chat participant is in possession of a Curve25519 key pair, with the public portion of this pair available to all other participants via an API (not part of this document). The new sender key is encrypted to *each* other participant by a key derived through ECDH [1]. Each recipient is able to extract and decrypt this embedded sender key of a participant. Each client is responsible of tracking these user-specific sender keys, identified uniquely by a tuple $(participant, keyID)$.

A message payload then is encrypted with this particular sender key and embedded within the message. Future follow-up messages only need to state the ID of the key previously transferred to indicate which encryption key was used to protect the message.

3.1.1 Sender Keys

Each sender key is 128 bit long for use with AES in CTR mode.

3.1.2 Key IDs

Key IDs are 32 bit long, need to be unique for each sender, and strictly monotonously increasing. For the purpose of the implementation they contain a 16 bit (high) portion derived from a UNIX Epoch time stamp to the granularity of days (incremented each UTC day), combined with a 16 bit (low) portion of a counter (starting from zero). The client implementation must prevent collisions or roll-overs.

3.1.3 Key Rotation

Whenever a new sender key is required, the sender will generate one, and send it (encrypted to all participants) along with the new key ID to all participating recipients in a keyed message. This is desirable to refresh a sender key (preventing extensively long use [2]), or when the composition of the group chat has changed (added and/or removed participants). Upon changes in the group composition the first message a client sends to the group chat *must* be a keyed message stating a new sender key.

[1] Note on older legacy clients: Some users may not have used a client recently that will generate a Curve25519 key pair. To not undermine the user experience, sender keys to such users will be encrypted with their RSA public keys. Upon logging into a client with a working Strongvelope implementation a Curve25519 key pair for the chat will be generated and used from then on.

[2] Our implementation rotates a sender key every 16 sent messages.

For convenience (e. g. when loading the chat history in reverse order), the previously used key with its key ID are re-sent to previous participants in the group as well. The client *must not* send the previous key to newly joined participants, and *must not* send a new key to departed participants.

3.1.4 Key Re-Sending

Sometimes the client needs to access messages out of context of a current message exchange flow, for example when loading previous messages from the chat history [3]. In those cases, a message's payload may only be decrypted with knowledge of the sender key used. Each message does carry the sender key ID used, but the (encrypted) key itself may be contained in one of the previous messages only. To ease the access to previous sender keys used, clients will re-send their current sender keys (including the respective key ID) in regular intervals. Ideally, this interval [4] is balanced with the intervals used for key rotation [2] and batch size for chat history loading [3]. A client may face the situation that key re-sending is due, but no chat message is to be sent. Therefore such a message can be sent without a user contributed payload resulting in a "blind message" (not to be displayed by the client application).

3.2 Content Encryption

Individual message content components need to be protected through encryption. When using a symmetric key several times, a different initialisation vectors (IVs) or nonces *must* be used. Each message carries one such "message nonce" encoded (but not encrypted) within the message.

3.2.1 Sender Key Encryption

Sender keys are encrypted using AES in CBC mode, using an initialisation vector (IV) along with the ECDH shared secret [1] with each participant.

To derive the shared secret one's own private (S_{own}) and the other participant's public key (S_{own}) is used. Through Curve25519 Diffie-Hellman scalar multiplication (ECDH) and subsequent application of a key derivation function (KDF, specifically *[HKDF]*-SHA256) the key is derived. It is trimmed to the required key size (128 most significant bits). According to RFC-5869 (*[HKDF]*) as a context info string the byte sequence "`strongvelope pairwise key`" is used, followed by the byte `0x01`.

$$K_{DH,dest} = KDF(ECDH(S_{own}, P_{other}))$$

To derive an IV (initialisation vector) for a recipient, the Mega user handle of the recipient is base64 URL decoded, yielding an 8 byte (64 bit) sequence (u). From these bytes then a keyed-hash message authentication code (HMAC, specifically HMAC-SHA-256) is computed using the message's master nonce (n) as a key, and subsequently trimmed to the required IV size (128 most significant bits).

$$IV_{dest} = HMAC(n_{master}, u_{other})$$

3.2.2 Payload Encryption

Message payloads are encrypted using AES in CTR mode, using the sender key and message nonce derived via computing an HMAC using the message's master nonce as a key and the byte sequence "`payload`" as a value. The message nonce will be trimmed to use the 96 most significant bits (12 bytes) only, leaving 32 bits for the counter.

$$n_{message} = HMAC(n_{master}, \text{"}\texttt{payload}\text{"})$$

[3] Our implementation loads chat history messages in batches of 32 messages.
[4] Our implementation re-sends sender keys every 30 total messages (received and sent).

CHAPTER

FOUR

PROTOCOL ENCODING

All fields in the messages exchanged are encoded as TLV (type, length, value) records. The entire message is prepended by a single byte indicating the protocol version (in case of future changes). At the moment the protocol version is `0x00`.

4.1 TLV Types

TLV records do not need to be in order according to the TLV type numbers, as long as it is assured that the `SIGNATURE` record is preceding all others.. Individual records may be missing or repeated multiple times. The currently defined TLV record types are:

Type `0x01` (1): `SIGNATURE`
Signature for all following bytes.

Type `0x02` (2): `MESSAGE_TYPE`
Type of message sent.

Type `0x03` (3): `NONCE`
"Message nonce" used for encryption (individual nonces are derived from it).

Type `0x04` (4): `RECIPIENT`
Recipient of message. This record is repeated for all recipients of the message.

Type `0x05` (5): `KEYS`
Message encryption keys, encrypted to a particular recipient. This may contain two (concatenated) keys. The second one (if present) is the previous sender key. Records of this type need to occur with the same number of records in the same order as `RECIPIENT`.

Type `0x06` (6): `KEY_IDS`
Sender encryption key IDs used (or set) in this message. If some `KEYS` records will contain a second (previous) key ID, the previous key ID needs to be given (concatenated) as well.

Type `0x07` (7): `PAYLOAD`
Encrypted payload of message.

Type `0x08` (8): `INC_PARTICIPANT`
Participant to be included with this message.

Type `0x09` (9): `EXC_PARTICIPANT`
Participant to be excluded with this message.

Type `0x0a` (10): `OWN_KEY`
Own message encryption (sender) key. This is usually not required, but will be used if legacy RSA encryption of sender keys is used for at least one recipient. This is required to access one's own sender key later when re-reading own chat messages later from history.

4.2 Message Signatures

All message signatures are detached cryptographic signature, signing the entire following (encoded) binary message content. To authenticate the entire message content, the signature *must* be contained in the first TLV record. All message signatures are computed using the sender's Ed25519 identity key.

The content to sign is computed as follows:

$$(magic\ number||content\ to\ sign)$$

Here, $magic\ number$ is a fixed string to distinguish the authenticator from any other content (for now it is the byte sequence "`strongvelopesig`").

4.3 Message Types

The message type is indicated by a single byte. Currently these types are in use:

Type `0x00` (0): `GROUP_KEYED`
Message containing a sender key (initial, key rotation, key re-send).

Type `0x01` (1): `GROUP_FOLLOWUP`
Message using an existing sender key for encryption.

Type `0x02` (2): `ALTER_PARTICIPANTS`
Alters the list of participants for the group chat (inclusion and exclusion).

4.4 Keyed Messages

After the mandatory byte for the protocol version, all keyed messages must contain the following records for a group chat of n participants:

- (0x01) `SIGNATURE`
- (0x02) `MESSAGE_TYPE`
- (0x03) `NONCE`
- (0x04) `RECIPIENT` [$(n-1)$ records of this type]
- (0x05) `KEYS` [$(n-1)$ records of this type, corresponding order to `RECIPIENT` records]
- (0x06) `KEY_IDS`
- (0x07) `PAYLOAD` [may only be omitted on "blind" messages, when re-sending keys]

4.5 Followup Messages

Followup messages are designed to be leaner (in terms of storage/transport size), and therefore are missing some of the TLV records. Messages contain the following TLVs (in the given order):

- (0x01) `SIGNATURE`
- (0x02) `MESSAGE_TYPE`
- (0x03) `NONCE`
- (0x06) `KEY_IDS` [contains only current/used key ID]
- (0x07) `PAYLOAD`

4.6 Alter Participant Messages

When altering the group composition of the chat (inclusion and/or exclusion of participants), the sender key must be rotated and the message must distribute the new sender key to all new participants. After the mandatory byte for the protocol version, all participant change messages must contain the following records for a group chat of n participants:

- (0x01) `SIGNATURE`
- (0x02) `MESSAGE_TYPE`
- (0x03) `NONCE`
- (0x04) `RECIPIENT` [$(n-1)$ records of this type]
- (0x05) `KEYS` [$(n-1)$ records of this type, corresponding order to `RECIPIENT` records]
- (0x06) `KEY_IDS`
- (0x08) `INC_PARTICIPANT` [a TLV record for each new participant to include]
- (0x09) `EXC_PARTICIPANT` [a TLV record for each former participant to exclude]
- (0x07) `PAYLOAD` [may only be omitted on "blind" messages, when re-sending keys]

4.7 Legacy Sender Key Encryption

In case at least one recipient of a keyed message (`GROUP_KEYED` or `ALTER_PARTICIPANTS` message types) requires encryption of a sender key using the recipient participant's RSA public key, the sender must add an additional TLV record of the `OWN_KEY` type. This `OWN_KEY` record is identical to a `KEYS` record, with the only difference that the key(s) are to be encrypted using one's own Curve25519 public key to derive a shared secret. This is required as one does not have access to a recipient's RSA private key to recover a sender key encrypted with it (e.g. when loading transport encrypted messages from the chat server's history).

4.8 Sender Keys to Include

When loading the history messages from a chat server on demand ("backwards" in history) one cannot easily decrypt preceding messages before a key rotation, until the next prior key rotation message has been retrieved as well. To ease this process, a client encrypts the new sender key to the recipient as well as the previous sender key (if possible and the recipient is entitled to it).

When including previous sender keys, also the `KEY_IDS` record must also include the previous sender key's ID (concatenated to the current sender key ID). Additionally the sending client *must* ensure that only entitled participants will receive this previous sender key. For example newly included participants are not entitled access to the content of previous messages, and therefore *must not* receive such previous sender keys.

CHAPTER

FIVE

REQUIREMENTS FOR MESSAGING WORKFLOWS

To assure proper operation of messaging workflows among participants, especially when encountering inclusions and/or exclusions of participants in the chat room, some common procedures must be ensured by the client.

Regarding sender keys to be included in keyed messages (`GROUP_KEYED` or `ALTER_PARTICIPANTS` message types) one *must* follow the requirements as outlined for alter participant messages (see *Alter Participant Messages*).

5.1 Initialising Chat Room Participation

When a client starts to participate in a chat session, two scenarios are possible: The chat session could be new for the client who has never participated in it, or the client could resume a chat it has previously participated in.

- In the former case (a new chat), the client must generate a new sender key and a corresponding key ID before sending the first message (see *Update Sender Key*). The first message then must then contain this sender key encrypted to all participants (keyed message).
- In the latter case (resuming a chat), the client must be "seeded" with chat history messages (see *Seed Encryption Handler*) to identify and extract own previous sender keys, as well as participants' sender keys along with their key IDs. In case the own client has never sent a message in the chat before, one reverts to the "new chat" behaviour outlined above (see also *Update Sender Key*).

5.2 Update Sender Key

When updating a sender key, a user's encryption handler generates a new key ID and associated sender key. With the next message sent, this new sender key and key ID are sent to the participants of the chat (a keyed message). If a previous sender key is available as well, it will also be included to those participants only who are entitled to it (participants in the chat when that sender key was in use). See *Alter Participant Messages* for the case of a changed chat group composition.

5.3 Seed Encryption Handler

When resuming a chat (starting to use a non-new chat, e.g. by responding to a previously sent initial chat message), the encryption handler needs to be "seeded" with sender keys. These include one's own sender keys and key IDs. For seeding, the handler will be given a suitably sized batch of messages from the message history to extract sender keys and their corresponding key IDs. The handler then will need to provide feedback on whether one's own (latest) sender key was contained in it. In the case it was *not*, a further (older, adjoining) batch of messages need to be passed to the seeder. In the case no own sender key could be found in the entire history available, the client must generate a new sender key for further use (see *Update Sender Key*).

A client may consider to "rotate" one's own sender key immediately at the beginning of a session (see *Rotate Key*).

5.4 Rotate Key

After a while of usage, a client may opt to "rotate" its sender key. That is, it generates a new one with a new key ID, and sends out a keyed message as the next message sent to the chat, informing all participants of the new sender key in use. This strategy is also used for the case of altering participants in a group chat. It ensures that previous members will be disallowed any access to newly sent messages, as well as new members are disallowed access to previous messages on a cryptographic level (see *Alter Participant Messages*).

5.5 Key Reminders

In situations with long one-sided communication, the history may be flooded with messages from one participant, and key rotation messages of oneself or other members may be scarce. This makes the seeding process (see *Seed Encryption Handler*) difficult, as large amounts of chat history may need to be fetched to successfully seed an encryption handler. Therefore, clients can send a keyed key reminder message in (regular) intervals. These messages may also be "blind" messages without content, just injecting the sender key material into the chat history again to aid future resumption.

5.6 Alter Participant Messages

The process triggered by alter participant messages is a bit more involved. In an alter participant message, one or more participants can included into or excluded from the group chat room. A client sending such a message will alter its local current group composition set, and include a list of participants to be included and excluded from the chat. Before sending such a message, the client rotates its sender key, and notifies the newly legitimate participants of the chat of this new sender key. The sender needs to ensure that the previous sender key *will not* be encrypted to newly included participants.

However, the recipients of an alter participant message also need to take action upon receiving it. They need to maintain a set for participants to include as well as to exclude on top of the current list of current participants. When they are to send their next own message, they need to rotate their sender key as well, and encrypt it to the now legitimate group chat participants similarly to the sender of the alter participant message. In the same fashion, included participants must not be informed of the previous sender key, and excluded participants must not receive any sender keys at all with that message.

Due to the fact that a number of alter participant messages may be received before the next own message is sent, the recipient needs to update their sets for participants to include and to exclude until such a first message is sent again. Once an own keyed message has been sent, the client may clear their include and exclude participant sets as well as their set of current group chat participants.

CHAPTER

SIX

STRONGVELOPE EVOLUTION

Strongvelope is not finished in its development/evolution, yet. The following are some notes on aspects still to be addressed.

- Concept of a “room operator” who has moderator privileges, and thus may be entitled for privileged operations (such as inclusion or exclusion of participants, setting of room topics/titles, etc.)
- ...

[mpENC] https://docs.mega.nz/chat/mpenc/mpENC.pdf

[Ed25519] http://ed25519.cr.yp.to/

[Curve25519] http://cr.yp.to/ecdh.html

[HKDF] https://tools.ietf.org/html/rfc5869

[NIST-ECC-Failure] Daniel J. Bernstein, Tanja Lange.
Failures in NIST's ECC standards.
https://cr.yp.to/newelliptic/nistecc-20160106.pdf